HP IO Accelerator 3.2.3 Solaris User Guide

Abstract

This document describes software requirements for all relevant HP IO Accelerators using Linux operating systems. This document is intended for system administrators who plan to install and use HP IO Accelerators with a Linux operating system. It is helpful to have previous experience with HP IO Accelerators and a Linux operating system. This user guide is intended for IO Accelerator software release 3.2.3 or later.

Part Number: 652148-001
March 2013
Edition: 1

© Copyright 2013 Hewlett-Packard Development Company, L.P.

The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

AMD is a trademark of Advanced Micro Devices, Inc.

Windows® is a U.S. registered trademark of Microsoft Corporation.

Contents

About this guide 5
Contents summary 5
Introduction 6
Overview 6
Product naming 6
Performance attributes 7
Required operating environment 8
Supported firmware revisions 8
Supported hardware 8
Software installation 11
Installation overview 11
Installing the software and utilities 11
Loading the IO Accelerator VSL facility (driver) 12
Upgrading the firmware 12
Enabling PCIe power 13
Using the device as a swap 13
Configuring a ZFS pool 13
Maintenance 14
Maintenance tools 14
Command-line utilities 14
Enabling PCIe power override 14
Enabling the override parameter 15
Uninstalling IO Accelerator VSL and utilities 16
Disabling auto attach 16
Unmanaged shutdown issues 16
Disabling the IO Accelerator VSL 17
Enabling the IO Accelerator VSL 17
Monitoring and managing devices 18
Introduction to monitoring and managing devices 18
Management tools 18
Example conditions to monitor 19
Device LED indicators 20
Performance and tuning 21
Introduction to performance and tuning 21
Disabling DVFS 21
Limiting APCI C-states 21
Utilities 22
Utilities reference 22
fio-attach 22
fio-beacon 23
fio-bugreport 23
fio-detach 24
fio-format 25

fio-status 26
fio-sure-erase 28
fio-update-iodrive 30
Monitoring IO Accelerator health 33
NAND flash and component failure 33
Health metrics 33
Health monitoring techniques 34
ZFS pools and health monitoring 35
Using module parameters 35
Upgrading devices 36
Upgrading devices for IO Accelerator VSL 2.x.x to 3.x.x 36
Upgrade procedure 36
Resources 39
Subscription service 39
For more information 39
Support and other resources 40
Before you contact HP 40
HP contact information 40
Customer Self Repair 40
Acronyms and abbreviations 48
Documentation feedback 49
Index 50

About this guide

Contents summary

- Instructions on downloading and installing the approved driver and utilities
- Instructions on maintaining the IO Accelerator
- Description of the following IO Accelerator models:
  - HP IO Accelerator for BladeSystem c-Class
  - HP PCIe IO Accelerator
  - HP PCIe IO Accelerator Duo

△ **CAUTION:** Before upgrading to 3.x.x software and firmware, back up all data on the IO Accelerator. The 3.2.3 software and firmware reformat the drive, which causes data to be lost if not backed up. The 3.2.3 software is not backward compatible with 1.2.x or 2.x software.

Introduction

Overview

Designed around ioMemory, a revolutionary storage architecture, HP IO Accelerator is an advanced NAND flash storage device. With performance comparable to DRAM and storage capacity on par with hard disks, the IO Accelerator increases performance so that every server can contain internal storage that exceeds the I/O performance of an enterprise SAN.

HP IO Accelerator is the first data accelerator designed specifically to improve the bandwidth for I/O-bound applications.

In addition to the hardware driver, the IO Accelerator also includes a VSL. This hybrid of the RAM virtualization subsystem and the disk I/O subsystem combines the best features of both systems. VSL functions as a disk to interface well with block-based applications and software, while also running like RAM underneath to maximize performance. This feature produces the following benefits:

- **Performance:** The VSL offers direct and parallel access to multiple CPU cores, enabling near linear performance scaling, consistent performance across different read/write workloads, and low latency with minimal interruptions and context switching.
- **Extensibility:** The VSL enables flash-optimized software development, making each IO Accelerator module a flexible building block for creating a flash-optimized data center.

Product naming

HP IO Accelerator Generation 1 devices include:

- AJ876A: HP 80 GB IO Accelerator for BladeSystem c-Class
- AJ877A: HP 160 GB IO Accelerator for BladeSystem c-Class
- AJ878A: HP 320 GB IO Accelerator for BladeSystem c-Class
- AJ878B: HP 320 GB IO MLC Accelerator for BladeSystem c-Class
- BK836A: HP 640 GB IO MLC Accelerator for BladeSystem c-Class

IMPORTANT: Generation 1 IO accelerators for BladeSystem c-Class are only compatible with G7 and earlier server blades.

- 600278-B21: HP 160 GB Single Level Cell PCIe ioDrive for ProLiant Servers
- 600279-B21: HP 320 GB Multi Level Cell PCIe ioDrive for ProLiant Servers
- 600281-B21: HP 320 GB Single Level Cell PCIe ioDrive Duo for ProLiant Servers
- 600282-B21: HP 640 GB Multi Level Cell PCIe ioDrive Duo for ProLiant Servers
- 641027-B21: HP 1.28 TB Multi Level Cell PCIe ioDrive Duo for ProLiant Servers

HP IO Accelerator Generation 2 devices include:

- QK761A: HP 365 GB IO MLC Accelerator for BladeSystem c-Class

- QK762A: HP 785 GB IO MLC Accelerator for BladeSystem c-Class
- QK763A: HP 1.2 TB IO MLC Accelerator for BladeSystem c-Class

IMPORTANT: Generation 2 IO accelerators for BladeSystem c-Class are only compatible with Gen8 and later server blades.

- 673642-B21: HP 365 GB Multi Level Cell G2 PCIe ioDrive2 for ProLiant Servers
- 673644-B21: HP 785 GB Multi Level Cell G2 PCIe ioDrive2 for ProLiant Servers
- 673646-B21: HP 1205 GB Multi Level Cell G2 PCIe ioDrive2 for ProLiant Servers
- 673648-B21: HP 2410 GB Multi Level Cell G2 PCIe ioDrive2 Duo for ProLiant Servers
- 721458-B21: HP 3.0 TB Multi Level Cell G2 PCIe ioDrive2 for ProLiant Servers

Performance attributes

IO Accelerator capacity Models AJ878B and BK836A	**320GB**	**640GB**
NAND type	MLC (Multi Level Cell)	MLC (Multi Level Cell)
Read Bandwidth (64kB)	735 MB/s	750 MB/s
Write Bandwidth (64kB)	510 MB/s	550 MB/s
Read IOPS (512 Byte)	100,000	93,000
Write IOPS (512 Byte)	141,000	145,000
Mixed IOPS* (75/25 r/w)	67,000	74,000
Access Latency (512 Byte)	30 µs	30 µs
Bus Interface	PCI-Express x4	PCI-Express Gen1 x4

IO Accelerator capacity Models QK762A and QK763A	**785GB**	**1.2TB**
NAND type	MLC (Multi Level Cell)	MLC (Multi Level Cell)
Read Bandwidth (1MB)	1.5 GB/s	1.5 GB/s
Write Bandwidth (1MB)	1.1 GB/s	1.3 GB/s
Read IOPS (Seq. 512 Byte)	443,000	443,000
Write IOPS (Seq. 512 Byte)	530,000	530,000
Read IOPS (Rand. 512 Byte)	141,000	143,000
Write IOPS (Rand. 512 Byte)	475,000	475,000
Read Access Latency	68µs	68µs
Write Access Latency	15µs	15µs
Bus Interface	PCI-Express Gen2 x4	PCI-Express Gen2 x4

*Performance achieved using multiprocessor enterprise server

- Enterprise data integrity
- Field upgradeability
- Green footprint, 7.5W nominal per device

NOTE: MSI was disabled to obtain these statistics.

Required operating environment

The HP IO Accelerator with software 3.2.3 is supported for use in the following operating environments:

- Red Hat Enterprise Linux 5 (AMD64/EM64T)
- Red Hat Enterprise Linux 6 (AMD64/EM64T)
- SUSE LINUX Enterprise Server 10 (AMD64/EM64T)
- SUSE LINUX Enterprise Server 11 (AMD64/EM64T)

CAUTION: Version 3.1.0 or greater of the driver software is not backward-compatible with any previous driver version. When you install version 3.2.3, you cannot revert to any previous version.

IMPORTANT: All operating systems must be 64-bit architecture.

NOTE: IO Accelerators cannot be used as hibernation devices.

Supported firmware revisions

After February 19, 2013, all IO Accelerators ship with firmware version 7.1.13.109322 or higher. This firmware version only works with VSL 3.2.2 or higher. If you are installing a recently purchased or a replacement IO Accelerator into a system that already has IO Accelerators installed, then you must upgrade the firmware on the previously installed devices to 7.1.13.109322 or higher. You must also upgrade the VSL to 3.2.2 or higher. Upgrading the firmware and VSL on cards that were running firmware versions 6.x.x or higher and VSL 3.x.x or higher is NOT data destructive. However, HP recommends that you back up any data on the device before performing the upgrade. The latest supported version of the firmware and VSL can be found on the HP website (http://www.hp.com).

Release	Firmware revision
1.2.4	17350
1.2.7	36867 or 42014
1.2.8.4	43246
2.2.x	43674
2.2.3	101583
2.3.1	101971 4 or 101971_6
3.1.1	107004 or greater
3.2.3	109322

Supported hardware

HP IO Accelerator for BladeSystem c-Class

BladeSystem c-Class IO Accelerators have two distinct designs for the respective server product lines. The G1 through G7 IO Accelerator adapter is provided in a c-Class Type 1 Mezzanine card form factor. It can be installed in both Type 1 and Type 2 mezzanine slots within the c-Class blade G1 through G7 servers, enabling a total of two cards in a half-height server blade, and three cards in a full-height server blade and up to 6 in a double-high, double-wide server (BL680c).

The Gen8 adapter is provided in a c-Class Type B Mezzanine card form factor. It can only be installed in Type B mezzanine slots within the Gen 8 or later servers, enabling one IO Accelerator in a half-height Gen8 server.

The Type I mezz card and the Type B mezz card are distinguished by the mezzanine connector. The type B card is slightly larger than a Type I mezz card.

The amount of free RAM required by the driver depends on the size of the blocks used when writing to the drive. The smaller the blocks, the more RAM is required. The table below lists the guidelines for each 80GB of storage. For the latest information, see the *QuickSpecs* sheet to the *HP IO Accelerator for HP BladeSystem c-Class* at HP Customer Support (http://www.hp.com/support).

The Remote Power Cut Module for the c-Class blade mezzanine card provides a higher level of protection in the event of a catastrophic power loss (for example, a user accidentally pulls the wrong server blade out of the slot). The Remote Power Cut Module ensures in-flight writes are completed to NAND flash in these catastrophic scenarios. Without the Remote Power Cut Module, write performance is slower. Writes are not acknowledged until the data is written to the NAND module, thereby slowing performance. When the Remote Power Cut Module is installed, writes are acknowledged by the IO Accelerator controller to the driver. The IO Accelerator controller then completes the write to the NAND module.

The IO Accelerators (QK761A, QK762A, and QK763A) for Gen 8 BladeSystem c-Class have the power cut functionality embedded on the card. They offer the protection without requiring the remote power cut module.

> **NOTE:** The Remote Power Cut Module is used only in the AJ878B and BK836A models. Without the Remote Power Cut Module, write performance is slower.

HP PCIe IO Accelerator minimum requirements

- An open PCI-Express slot—The accelerator requires a minimum of one half-length, half-height slot with a x4 physical connector. All four lanes must be connected for full performance. HP PCIe IO Accelerator Duo requires a minimum of a full-height, half-length slot with an x8 physical connection. If your system is using PCI 1.1, all x8 signaling lanes must be connected for full performance. If your system is using PCI 2.0, for full performance you only have to connect x4 signaling lanes.

  > **NOTE:** For PCIe IO Accelerators, using PCIe slots greater than x4 does not improve performance.

  > **NOTE:** The power cut feature is built into PCIe IO Accelerators; therefore, no Remote Power Cut Module is necessary.

- 300 LFM of airflow at no greater than 50°C. To protect against thermal damage, the IO Accelerator also monitors the junction temperature of its controller. The temperature represents the internal temperature of the controller, and it is reported in `fio-status` report. The IO Accelerator begins throttling write performance when the junction temperature reaches 78°C. If the junction temperature continues to rise, the IO Accelerator shuts down when the temperature reaches 85°C.

  > **NOTE:** If you experience write performance throttling due to high temperatures, see your computer documentation for details on increasing airflow, including fan speeds.

- Sufficient RAM to operate—The amount of RAM that the driver requires to manage the NAND flash varies according to the block size you select when formatting the device (filesystem format, not low-level format). For a virtual machine using an IO Accelerator directly (using PCI pass-through), consult the user guide for the installed operating system. The following table lists the amount of RAM required per 100GB of storage space, using various block sizes. The amount of RAM used in driver version 3.0 is significantly less than the amount used in version 1.2.x.

Average block size (bytes)	RAM usage for each 80 GB IO Accelerator (Megabytes)	RAM usage for each 100 GB IO Accelerator (Megabytes)	Minimum system RAM requirement for 320 GB MezzIO Accelerator*
8,192	250	280	1 GB
4,096	400	530	1.6 GB
2,048	750	1,030	3 GB
1,024	1,450	2,000	5.8 GB
512	2,850	3,970	11.4 GB

Average block size (bytes)	Minimum system RAM requirement for 640 GB MezzIO Accelerator*	Minimum system RAM requirement for 785 GB MezzIO Accelerator*	Minimum system RAM requirement for 1.2 TB MezzIO Accelerator*
8,192	2 GB	2.2 GB	3.4 GB
4,096	3.2 GB	4.2 GB	6.4 GB
2,048	6 GB	8.1 GB	12.4 GB
1,024	11.6 GB	15.7 GB	24 GB
512	22.8 GB	31.2 GB	47.6 GB

* For IO Accelerator use only. Additional RAM is needed for system operating system and applications.

HP PCIe IO Accelerator Duo requirements

In addition to the IO Accelerator cooling and RAM requirements listed in the previous table, the IO Accelerator Duo requires at least:

- A PCIe Gen1 x8 slot or a PCIe Gen2 x4 slot
- A minimum of a full-height, half-length slot with a x8 physical connection. For systems with PCI 1.1, all eight signaling lanes must be active for full IO Accelerator Duo performance. For systems with PCIe 2.0, only four signaling lanes must be active for full performance.

  NOTE: With driver version 3.1.0 and later, the driver detects in the BIOS if the PCIe slot supports a 75W power draw. If the slot supports up to 75W, the IO Accelerator device draws up to that amount of power. However, if an external power cable is used, power is only supplied by that cable.
  To verify whether a slot is supplying 75W, view the system logs or use the `fio-pci-check` utility.

Software installation

Installation overview

Before continuing with the installation of this software, do the following:

1. Ensure that your operating system is included in the list of supported operating systems contained in the release notes for this release.
2. Before installing the IO Accelerator VSL, make sure you have properly installed the IO Accelerator devices. For more information, see the IO Accelerator hardware installation guide. Every IO Accelerator device in a system must be upgraded to the appropriate firmware requirements.

Multiple devices in a system

Every IO Accelerator device in a system must be upgraded to the appropriate firmware.

For example, if you have a system running IO Accelerator software version 2.3.1 with IO Accelerator devices previously installed and want to install new IO Accelerator devices that require the latest version of the firmware, then you must upgrade all of the existing devices with firmware that supports this version of the IO Accelerator VSL software.

The IO Accelerator VSL for Solaris is distributed as a tarball.

Installing the software and utilities

To install the IO Accelerator VSL software and utilities:

1. Download the appropriate tarball from the HP website (http://www.hp.com/go/support).

   Example:
   `hp_io_accelerator-<version>.pkg.tar.gz`
2. Change to the directory containing the downloaded tarball.
3. Uncompress the `-vsl` package:
   ```
   $ gunzip -c iomemory-vsl*pkg.tar.gz|tar -xf -
   ```
4. Run the `pkgadd` command:
   ```
   $ pfexec pkgadd -d . iomemory-vsl
   ```

The software and utilities are installed to the following locations.

File type	Installation location
Driver	`/usr/kernel/drv/amd64`
Utilities	`/opt/fusionio/bin`
Driver Config File	`/usr/kernel/drv`

Continue to "Loading the IO Accelerator VSL facility (driver) (on page 12)."

Loading the IO Accelerator VSL facility (driver)

The driver is installed and automatically loaded during the package installation (`pkgadd`; see "Installing the software and utilities (on page 11)").

To manually unload the driver, run:

```
$ pfexec rem_drv iomemory-vsl
```

To manually load the driver, run:

```
$ pfexec add_drv -i "pci1aed,1005 pciex1aed,1005 pci1aed,1003 pciex1aed,1003
pci1aed,1010 pciex1aed,1010 pci1014,3c3 pci103c,324d pci103c,324e
pci103c,178b pci103c,178c pci103c,178d pci103c,178e pci103c,176f
pci103c,1770" -m "* 0666 root sys" iomemory-vsl
```

The load command might be too long for the shell. You might need to add a \ (backslash) character at the end of the first line to signal that the line is continued on the next line.

The driver automatically loads at system boot. The IO Accelerator device is now available to the OS as the following:

```
/dev/rdsk/c*d0p0
```

Attached IO Accelerator devices show up as block devices in the file system in three places:

```
/dev/fio*
/dev/rdsk/c*d0*
```

Where `c*d0p0` is the master partition covering the entire raw block device

```
/dev/dsk/c*d0*
```

Where `c*d0s2` is the master slice covering the entire block device

The driver creates `/dev/fiox` block device nodes, where `x` is a letter (a, b, c, and so forth).

> **CAUTION:** Only use the traditional disk interfaces (`/dev/rdsk/*` and `/dev/dsk/*`). Do not use the `/dev/fiox` block device nodes.

To confirm the IO Accelerator device is attached, use the following command:

```
$ pfexec /opt/fusionio/bin/fio-status
```

The output indicates the attach status of the device: attached or not attached.

Upgrading the firmware

After the IO Accelerator driver is loaded, ensure that the IO Accelerator device firmware is up to date by running the "`fio-status (on page 26)`" command-line utility.

If the output shows that the device is running in Minimal mode, download the latest firmware from the HP website (http://www.hp.com/support), and then use the HP IO Accelerator Management Tool application or the "`fio-update-iodrive (on page 30)`" utility to upgrade the firmware.

> **CAUTION:** Upgrade Path
>
> - Do not attempt to downgrade the firmware on any IO Accelerator device.
> - You must follow a specific upgrade path when upgrading an IO Accelerator device.
> - When installing a new IO Accelerator device along with existing devices, you must upgrade all of the existing devices to the latest available versions of the firmware before installing the new devices.
> - Consult the release notes for this IO Accelerator release before upgrading IO Accelerator devices.

IMPORTANT: The IO Accelerator device might have a minimum firmware label affixed (for example, "MIN FW: XXXXXX"). This label indicates the minimum version of the firmware that is compatible with your device.

Enabling PCIe power

For PCIe IO Accelerators, if you have installed any dual IO Accelerator devices, such as the HP ioDrive2 Duo, then the device might require additional power than the minimum 25 W provided by PCIe Gen2 slots to properly function.

For instructions on enabling the device to draw additional power from the PCIe slots, see "Enabling PCIe power override (on page 14)."

Using the device as a swap

In Solaris, HP does not recommend using the IO Accelerator device as a swap space because deadlocked memory issues might occur.

Configuring a ZFS pool

You can configure two or more IO Accelerator devices into a RAID array using a single ZFS pool.

If you are using RAID1/Mirrored and one device fails, be sure to run `fio-format` on the replacement device, not the existing good device, before rebuilding the RAID. The following examples of some common RAID configurations use ZFS.

In the RAID configurations in this section, `c12d0p0`, `c13d0p0`, `c14d0p0`, and `c15d0p0` are used as examples; they do not necessarily match your RAID installation.

Using a single ioMemory device

```
$ pfexec zpool create mypool c12d0p0
```

RAID 0

To create a striped set, where `c12d0p0` and `c13d0p0` are the two IO Accelerator devices you want to stripe, run the following command:

```
$ pfexec zpool create mypool c12d0p0 c13d0p0
```

RAID 1

To create a mirrored set using the two IO Accelerator devices `c12d0p0` and `c13d0p0`, run the following command:

```
$ pfexec zpool create mypool mirror c12d0p0 c13d0p0
```

RAID 10

To create a striped, mirrored array using four IO Accelerator devices, run the following command:

```
$ pfexec zpool create mypool mirror c12d0p0 c13d0p0 mirror c14d0p0 c15d0p0
```

Maintenance

Maintenance tools

The IO Accelerator includes software utilities for maintaining the device. You can also install SNMP as a monitoring option.

The following are the most common tasks for maintaining your IO Accelerator. You can also use the IO Accelerator Management Tool application to perform firmware upgrades. For more information, see the *HP IO Accelerator Management Tool User Guide*.

Command-line utilities

Several command-line utilities are included in the installation packages for managing your IO Accelerator device:

- `fio-attach`
- `fio-beacon`
- `fio-bugreport`
- `fio-detach`
- `fio-format`
- `fio-status`
- `fio-update-iodrive`

For more information on command-line utilities, see "Utilities (on page 22)."

Enabling PCIe power override

For PCIe IO Accelerators, if you have installed any dual IO Accelerator devices, such as the HP ioDrive2 Duo, then the device might require additional power than the minimum 25 W provided by PCIe Gen2 slots to properly function. Even if additional power is not required for your device, all dual IO Accelerator devices that receive additional power might benefit with improved performance.

HP ioDrive2 Duo devices must have additional power to properly function. For more information on which devices require additional power, see the *HP PCIe IO Accelerator for ProLiant Servers Installation Guide*.

Additional power can be provided in two ways:

- **External power cable**—This cable ships with all dual ioMemory devices. For information on installing this cable, see the *HP PCIe IO Accelerator for ProLiant Servers Installation Guide*.

  NOTE: When a power cable is used, all of the power is drawn from the cable and no power is drawn from the PCIe slot.

- **Enabling full slot power draw**—Some PCIe slots provide additional power (often up to 75 W of power). If your PCIe slot is rated to provide at least 55 W, then you can allow the device to draw full power from the PCIe slot by setting a VSL module parameter. For more information on enabling this override parameter, see "Enabling the override parameter (on page 15)."

  CAUTION: If the PCIe slot is not capable of providing the needed amount of power, then enabling full power draw from the PCIe slot might result in malfunction or even damage to server hardware. The user is responsible for any damage to equipment due to improper use of the override parameter. HP expressly disclaims any liability for damage arising from improper use of the override parameter. To confirm the power limits and capabilities of each slot, as well as the entire system, contact the server manufacturer. For information about using the override parameter, contact HP Customer Support.

  NOTE: The override parameter overrides the setting that prevents devices from drawing more than 25 W from the PCIe slot. The parameter is enabled by device using the device serial numbers. Once the setting is overridden, each device might draw up to the full 55 W needed for peak performance.

  Before you enable the override parameter, ensure that each PCIe slot is rated to provide enough power for all slots, devices, and server accessories. To determine the power slot limits, consult the server documentation, BIOS interface, setup utility, or use the `fio-pci-check` command.

Important considerations

- If you are installing more than one dual IO Accelerator device and enabling the override parameter for each device, be sure the motherboard is rated to provide 55W power to each slot used. For example, some motherboards safely provide up to 75W to any one slot, but run into power constraints when multiple slots are used to provide that much power. Installing multiple devices in this situation might result in server hardware damage. Consult with the manufacturer to determine the total PCIe slot power available.
- The override parameter persists in the system and enables full power draw on an enabled device even if the device is removed and then placed in a different slot within the same system. If the device is placed in a slot that is not rated to provide 55W of power, your server hardware could experience a power drag.
- The override parameter is a setting for the IO Accelerator VSL software by server and is not stored in the device. When moved to a new server, the device defaults to the 25 W power limit until an external power cable is added or the override parameter is enabled for that device in the new server. To determine the total PCIe slot power available for the new server, consult the manufacturer.

Enabling the override parameter

1. Use one of the following methods to determine the serial number of each device to be installed in a compatible slot:
   - Enter the `fio-status` command:

     Sample output:

     ```
     fio-status
     ...
     Adapter: Dual Controller Adapter
     Fusion-io ioDrive2 DUO 2.41TB, Product Number:F01-001-2T41-CS-0001, FIO
     SN:1149D0969
     ```

```
External Power: NOT connected
PCIe Power limit threshold: 24.75W
Connected ioMemory modules:
fct2: SN:1149D0969-1121
fct3: SN:1149D0969-1111
```

In this example, `1149D0969` is the adapter serial number.

If you have multiple IO Accelerator devices installed on your system, use the fio-beacon command to verify where each device is physically located.

- Inspect the adapter serial number labels on the IO Accelerator devices to determine the serial numbers. However, HP recommends confirming that each serial number is an adapter serial number by running the fio-status command. The adapter serial number label resides on the back of all HP ioDrive Duo devices and HP ioDrive2 Duo devices. On ioDrive Duos, the serial number is located on the PCB component that is attached to the PCIe connector.

2. To set the module parameter, edit the `/usr/modprove.d/iomemory-vsl.conf` file, and then change the value for the `external_power_override` parameter. For example:

   ```
   options iomemory-vsl external_power_override=<value>
   ```

   Where the <value> for this parameter is a comma-separated list of adapter serial numbers. For example:

   1149D0969,1159E0972,24589

3. To enforce any parameter changes, you must reboot or unload and then load the drivers.

Uninstalling IO Accelerator VSL and utilities

To uninstall the IO Accelerator VSL and utilities, run the following command:

```
$ pfexec pkgrm iomemory-vsl
```

Disabling auto attach

When the IO Accelerator driver is installed, it is configured to attach any devices automatically when the driver is loaded. When necessary, disable the auto-attach feature.

To disable auto attach:

1. Edit the following file:
   `/usr/kernel/drv/iomemory-vsl.conf`
2. Add the following line to that file:
   `auto_attach=0;`
3. Save the file.

To re-enable auto-attach, edit the file, and then either remove that line or change it to:

```
auto_attach=1;
```

Unmanaged shutdown issues

If you experience an unmanaged shutdown, then the IO Accelerator performs a consistency check during the reboot. The reboot might take several minutes or more to complete and is indicated by a progress percentage during the startup.

Although data written to the IO Accelerator device is not lost due to unmanaged shutdowns, important data structures might not have been properly committed to the device. This consistency check repairs these data structures.

Disabling the IO Accelerator VSL

The driver automatically loads by default when the operating system starts. You can disable IO Accelerator auto-load for diagnostic or troubleshooting purposes.

To disable IO Accelerator driver auto-load:

1. Use the following command:
   ```
   $ pfexec rem_drv iomemo
   ```
   The ioDrive ioMemory VSL does not load, so the device is not available to users.
2. Proceed with troubleshooting to correct the problem.

Enabling the IO Accelerator VSL

To re-enable the IO Accelerator VSL after maintenance, run the following command:

```
$ pfexec /usr/sbin/add_drv -i "pci1aed,1005" "pciex1aed,1005"
"pci1aed,1003"
"pciex1aed,1003" "pci1aed,1010" "pciex1aed,1010" "pci1014,3c3"
"pci103c,324d" "pci103c,324e" "pci103c,178b" "pci103c,178c" "pci103c,178d"
"pci103c,178e" "pci103c,176f" "pci103c,1770" -m "*0666 root sys"
iomemory-vsl
```

Monitoring and managing devices

Introduction to monitoring and managing devices

HP provides many tools for managing your IO Accelerator devices. These tools allow you to monitor the devices for errors, warnings, and potential problems. They also allow you to manage the devices, including performing the following functions:

- Firmware upgrades
- Low-level formatting
- Attach and detach actions
- Device status and performance information
- Configuring Swap and Paging
- Generating bug reports

Management tools

HP has provided tools for monitoring and managing IO Accelerator devices. These include stand-alone tools that require no additional software and a data-source tool that can be integrated with other applications.

Consider the descriptions of each tool to decide which tool (or combination of tools) best fits your needs.

About error messages

The IO Accelerator VSL software prints some error messages to the system logs. While these messages are useful for troubleshooting purposes, the VSL log messages are not designed for continual monitoring purposes (as each is based on a variety of factors that could produce different log messages depending on environment and use case). For best results, use the tools described in this section to regularly monitor your devices.

Stand-alone tool

These stand-alone tools do not require any additional software.

Command-line utilities: These utilities are run manually in the terminal. The `fio-status` utility provides status for all devices within a host. The other utilities allow you to perform other management functions. For more information, see "Utilities (on page 22)."

iData-source tool

This data-source tool provides comprehensive data, like the stand-alone tools, but it requires integration with additional software. At a minimum, it can interface with a browser. However, the benefit of this tool is that it can be integrated into existing management software that is customized for your organization.

This tool is available as a separate download. For more information, see the IO Accelerator release notes.

SNMP Subagent: The SNMP AgentX subagent allows you to monitor and manage your ioMemory devices using the Simple Network Management Protocol. You can use a normal SNMP browser or customize your existing application to interface with the subagent.

Example conditions to monitor

This section provides examples of conditions you can monitor. It is intended as an introduction and not as a comprehensive reference. These conditions have slightly different names, states, and values, depending on the tool you choose. For example, an SNMP MIB might have a different name than a SMI-S object or an API function.

To properly monitor these conditions, you should become familiar with the tool you choose to implement and read the documentation for that tool.

The possible states and values of these conditions are:

- GREEN: Normal
- YELLOW: Caution/Alert
- RED: Error/Warning

You can implement your own ranges of acceptable states/values, especially if you use a data-source tool.

Device status

All of the monitoring tools return information on the status of the IO Accelerator devices, including the following states:

- GREEN: Attached
- YELLOW: Detached, Busy (including: Detaching, Attaching, Scanning, Formatting, and Updating)
- RED: Minimal Mode, Power loss Protect Disabled

If the device is in Minimal Mode, the monitoring tool can display the reason for the Minimal Mode status.

Required actions

If the device is in Minimal Mode, the action depends on the reason. For example, if the reason is outdated firmware, then you must update the firmware.

Temperature

IO Accelerator devices require adequate cooling. To prevent thermal damage, the IO Accelerator VSL starts throttling write performance after the on-board controller reaches a specified temperature. If the controller temperature continues to rise, the software shuts down the device after the controller temperature reaches the maximum operating temperature.

These temperatures depend on the device. Newer IO Accelerator devices have higher thermal tolerances. Consult the IO Accelerator hardware installation guide to determine the thermal tolerances of all devices you will monitor. This list uses the thermal tolerances for newer devices (93°C throttling, 100°C shutdown).

- GREEN: <93°C
- YELLOW: 93-99°C
- RED: 100°C

You can shift the conditions by a few degrees so the YELLOW condition exists before throttling occurs. For example:

- GREEN: <90°C
- YELLOW: 90-96°C
- RED: 97°C

Required actions

If the temperature is at or approaching the YELLOW condition, you must increase the cooling for your system. This might include increasing the fan speed, bringing down the ambient temperature, reducing write load, or moving the device to a different slot.

Health reserves percentage

IO Accelerator devices are highly fault-tolerant storage subsystems with many levels of protection against component failure and the loss nature of solid-state storage. As in all storage subsystems, component failures can occur.

By proactively monitoring device age and health, you can ensure reliable performance over the intended product life. The following describes the Health Reserve conditions.

- GREEN: >10%
- YELLOW: 0-10%
- RED: 0%

At the 10% healthy threshold, a one-time warning is issued. At 0%, the device is considered unhealthy. It enters write-reduced mode. After the 0% threshold, the device enters read-only mode.

For complete information on Health Reserve conditions and their impact on performance, see "Monitoring IO Accelerator health (on page 33)."

Required actions

The device needs close monitoring as it approaches 0% reserves and goes into write-reduced mode, which results in reduced write performance. Prepare to replace the device soon.

Flashback protection

Like many other memory devices, NAND flash eventually fails with use. Those failures can be either permanent or temporary. Flashback redundancy is designed to address those chips that experience permanent failures, and provides additional protection above and beyond ECC for soft failures.

Flashback provides a real-time RAID-like redundancy at the chip-level, without sacrificing user capacity or performance for fault tolerance.

- GREEN: No Flashback indication
- YELLOW: Flashback indication
- RED: Failed Flashback indication

Required actions

The device functions normally in Flashback mode. Continue to monitor the device.

As a best practice, always back up your data on a regular basis. Flashback protection mode does not signal impending failure, but it is a reminder that devices can fail, and that data is best protected with proper redundancy.

If Flashback protection has failed, then the device is no longer usable. The device might need to be replaced. Run the `fio-bugreport` utility and contact HP Customer Support.

Device LED indicators

The IO Accelerator device includes three LEDs showing drive activity or error conditions.

Performance and tuning

Introduction to performance and tuning

HP IO Accelerator devices provide high bandwidth and high IOPS and are specifically designed to achieve low latency.

As IO Accelerator devices improve in IOPS and low latency, the device performance may be limited by operating system settings and BIOS configuration. To take advantage of the revolutionary performance of IO Accelerator devices, you might have to tune these settings.

While IO Accelerator devices generally perform well out of the box, this section describes some of the common areas where tuning may help achieve optimal performance.

Disabling DVFS

DVFS is a power management technique that adjusts the CPU voltage and frequency to reduce power consumption by the CPU. These techniques help conserve power and reduce the heat generated by the CPU, but they adversely affect performance while the CPU transitions between low-power and high-performance states.

These power-savings techniques are known to have a negative impact on I/O latency and maximum IOPS. When tuning for maximum performance, you might benefit from reducing or disabling DVSF completely, even though this might increase power consumption.

DVFS, if available, should be configurable as part of your operating systems power management features as well as within your system BIOS interface. Within the operating system and BIOS, DVFS features are often found under the ACPI sections. Consult your computer documentation for details.

Limiting APCI C-states

Newer processors have the ability to go into lower power modes when they are not fully utilized. These idle states are known as ACPI C-states. The C0 state is the normal, full power, operating state. Higher C-states (C1, C2, C3, and so on) are lower power states.

While ACPI C-states save on power, they are known to have a negative impact on I/O latency and maximum IOPS. With each higher C-state, typically more processor functions are limited to save power, and it takes time to restore the processor to the C0 state.

These power savings techniques are known to have a negative impact on I/O latency and maximum IOPS. When tuning for maximum performance you might benefit from limiting the C-states or turning them off completely, even though this might increase power consumption.

If your processor has ACPI C-states available, you can typically limit or disable them in the BIOS interface (sometimes referred to as a Setup Utility). APCI C-states might be part of the ACPI menu. For details, see your computer documentation.

Utilities

Utilities reference

The IO Accelerator installation packages include various command line utilities, installed by default to the `/usr/bin` file. These utilities provide a number of ways to access, test, and manipulate your device.

Utility	Purpose
fio-attach	Makes an IO Accelerator available to the OS
fio-beacon	Lights the IO Accelerator external LEDs
fio-bugreport	Prepares a detailed report for use in troubleshooting issues
fio-detach	Temporarily removes an IO Accelerator from OS access
fio-format	Used to perform a low-level format of an IO Accelerator
fio-status	Displays information about the device
fio-update-iodrive	Updates the IO Accelerator firmware

NOTE: All utilities have –h (Help) and –v (Version) options.

To set IO Accelerator VSL parameters, edit the `iomemory-vsl.conf` file:
`(/usr/kernel/drv/iomemory-vsl.conf)`

fio-attach

Description

Attaches the IO Accelerator device and makes it available to the operating system. This creates a block device in `/dev` named fiox (where x is a, b, c, and so on). Then, you might partition or format the IO Accelerator device, or set it up as part of a RAID array. The command displays a progress bar and percentage as it operates.

NOTE: In most cases, the IO Accelerator automatically attaches the device on load and does a scan. You only have to run `fio-attach` if you ran `fio-detach` or if you set the IO Accelerator `auto_attach` parameter to `0`.

Syntax

```
fio-attach <device> [options]
```

where `<device>` is the name of the device node (`/dev/fctx`), where `x` indicates the card number: `0`, `1`, `2`, and so on. For example, `/dev/fct0` indicates the first IO Accelerator device installed on the system.

You can specify multiple IO Accelerator devices. For example, `/dev/fct1 /dev/fct2` indicates the second and third IO Accelerator devices installed on the system. You can also use a wildcard to indicate all IO Accelerator devices on the system.

For example, `/dev/fct*`

Option	Description
-c	Attach only if clean.
-q	Quiet: disables the display of the progress bar and percentage.

fio-beacon

Description

The fio-beacon utility enables all three LEDs on, to identify the specified IO Accelerator device.

IMPORTANT: This utility enables the LEDs on, unless you select the -0 option.

Syntax

```
fio-beacon <device> [options]
```

where <device> is the name of the device node (/dev/fctx), where x indicates the card number: 0, 1, 2, and so on. For example, /dev/fct0 indicates the first IO Accelerator device installed on the system.

Options	Description
-0	Off: Turns off the three LEDs
-1	On (default): Enables the three LEDs on
-p	Prints the PCI bus ID of the device at <device> to standard output. Usage and error information may be written to standard output rather than to standard error.

fio-bugreport

Description

Prepares a detailed report of the device for use in troubleshooting problems. The results are saved in the /tmp directory in a file that indicates the date and time the utility was run.

Example

```
/tmp/fio-bugreport-20100121.173256-sdv9ko.tar.bz2
```

Syntax

```
fio-bugreport
```

NOTE: If the utility recommends that you contact Fusion-io support, disregard that message and contact HP support (http://www.hp.com/support) instead.

Sample output

```
-bash-3.2# fio-bugreport /tmp/fio-bugreport-20090921.173256-sdv9ko ~
Collecting fio-status -a
Collecting fio-status
Collecting fio-pci-check
Collecting fio-pci-check -v
```

```
Collecting fio-read-lebmap /dev/fct0
Collecting fio-read-lebmap -x /dev/stdout/dev/fct0
Collecting fio-read-lebmap -t /dev/fct0
Collecting fio-get-erase-count/dev/fct0
Collecting fio-get-erase-count -b /dev/fct0
Collecting lspci
Collecting lspci -vvvvv
Collecting lspci -tv
Collecting messages file(s)
Collecting procfusion file(s)
Collecting lsmod
Collecting uname -a
Collecting hostname
Collecting sar -r
Collecting sar
Collecting sar -A
Collecting syslog file(s)
Collecting proc file(s)
Collecting procirq file(s)
Collecting dmidecode
Collecting rpm -qa iodrive*
Collecting find /lib/modules
Please send the file /tmp/fio-bugreport-20090921.173256-sdv9ko.tar.bz2
along with your bug report to support@fusionio.com The file is in the /tmp
directory.
```

For example, the filename for a bug report file named

`/tmp/fiobugreport-20090921.173256-sdv9k0.tar.bz2` indicates the following:

- Date (20090921)
- Time (173256, or 17:32:56)
- Misc. information (sdv9ko.tar.bz2)

fio-detach

Description

Detaches and removes the corresponding `/dev/fctx` IO Accelerator block device. The `fio-detach` command waits until the device completes all read/write activity before executing the detach process. The command displays a progress bar and percentage as it completes the process.

> **NOTE:** Before using this utility, be sure that the device you want to detach is not currently mounted and in use.

Syntax

```
fio-detach <device> [options]
```

where <device> is the name of the device node (/dev/fctx), where x indicates the board number: 0, 1, 2, and so on. For example, /dev/fct0 indicates the first IO Accelerator installed on the system.

You can specify multiple IO Accelerator devices. For example, /dev/fct1 /dev/fct2 indicates the second and third IO Accelerator devices installed on the system. You can also use a wildcard to indicate all IO Accelerator devices on the system.

For example, /dev/fct*

Options	Description
-i	Immediate: Causes a forced immediate detach (does not save metadata). This will fail if the device is in use by the OS.
-q	Quiet: Disables the display of the progress bar and percentage.

NOTE: Detaching an IO Accelerator device might fail with an error indicating that the device is busy. This might occur if the IO Accelerator device is part of a software RAID volume (0,1,5), is mounted, or if some process has the device open.
The tools fuser, mount, and lsof can be helpful to determine what is holding the device open.

fio-format

Description

IMPORTANT: The IO Accelerator devices are shipped pre-formatted. fio-format is not required except to change the logical size or block size of a device, or to erase user data on a device. To ensure the user data is truly erased, use fio-sure-erase.

The fio-format utility performs a low-level format of the board. By default, fio-format displays a progress-percentage indicator as it runs.

CAUTION: Use this utility with care since it deletes all user information on the IO Accelerator.

IMPORTANT: Use the -s or -o option to change the default capacity of the device size. When enabled, the -s and -o options must include the size or percentage indicators.

NOTE: Use a large block (sector) size to reduce IO Accelerator memory consumption. For example: 4096 bytes. Be aware that some applications are not compatible with non-512-byte sector sizes.

Syntax

```
fio-format [options] <device>
```

where <device> is the name of the device node (/dev/fctx), where x indicates the device number: 0, 1, 2, and so on. For example, /dev/fct0 indicates the first IO Accelerator device installed on the system.

Options	Description
-b <size B|K>	Set the block (sector) size, in bytes or KiBytes (base 2). The default is 512 bytes. For example: -b 512B or -b 4K (B in 512B is optional).

Options	Description
-f	Force the format size, bypassing normal checks and warnings. This option may be needed in rare situations when fio-format does not proceed properly. (The "Are you sure?" prompt still appears unless you use the -y option.) This option can only be used with the -o option.
-q	Quiet mode: Disable the display of the progress-percentage indicator.
-s <size B|K|M|G|T|%>	Set the device capacity as a specific size (in TB, GB, or MB) or as a percentage of the advertised capacity: • T Number of terabytes (TB) to format • G Number of gigabytes (GB) to format • M Number of megabytes (MB) to format • % Percentage, such as 70% (The percent sign must be included)
-o <size B|K|M|G|T|%>	Over-format the device size (to greater than the preset capacity), where the maximum size equals the maximum physical capacity. If a percentage is used, it corresponds to the maximum physical capacity of the device. (Size is required for the -o option; see the -s option above for size indicator descriptions.) Before using this option, for supported recommendations, contact HP Support (http://www.hp.com/support)
-r	Enable fast rescan on non-ordered shutdowns at the cost of some capacity.
-y	Auto-answer "yes" to all queries from the application (bypass prompts).

fio-status

Description

Provides detailed information on installed devices. This utility operates on either fctx or fiox devices. The utility runs from root level and having the IO Accelerator loaded. If no IO Accelerator is loaded, there is less content reported from the query.

Enabling fio-status provides alerts for certain error modes, such as a minimal-mode, read-only mode, and write-reduced mode, describing what is causing the condition.

Syntax

```
fio-status [<device>] [<options>]
```

where <device> is the name of the device node (/dev/fctx), where x indicates the card number: 0, 1, 2, and so on. For example, /dev/fct0 indicates the first IO Accelerator device installed on the system.

If <dev> is not specified, fio-status displays information for all cards in the system. If the IO Accelerator is not loaded, this parameter is ignored.

Options	Description
-a	Report all available information for each device.
-e	Show all errors and warnings for each device. This option is for diagnosing issues, and it hides other information such as format sizes.
-c	Count: Report only the number of IO Accelerator devices installed.
-d	Show basic information set plus the total amount of data read and written (lifetime data volumes). This option is not necessary when the -a option is used.
-fj	Format JSON: creates the output in JSON format.
-fx	Format XML: creates the output in XML format.

Options	Description
`-u`	Show unavailable fields. Only valid with `-fj` or `-fx`.
`-U`	Show unavailable fields and details why. Only valid with `-fj` or `-fx`.
`-F<field>`	Print the value for a single field (For field names, see the next option). Requires that a device be specified. Multiple `-F` options may be specified.
`-l`	List the fields that can be individually accessed with `-F`.

△ **CAUTION: Output Change**
Starting with version 3.0.0 and later, the standard formatting of `fio-status` output has changed. This will affect any custom management tools that used the output of this utility.

Basic information: If no options are used, `fio-status` reports the following basic information:

- Number and type of devices installed in the system
- IO Accelerator version

Adapter information:

- Adapter type
- Product number
- External power status
- PCIe power limit threshold (if available)
- Connected IO Accelerator devices

Block device information:

- Attach status
- Product name
- Product number
- Serial number
- PCIe address and slot
- Firmware version
- Size of the device, out of total capacity
- Internal temperature (average and maximum, since IO Accelerator load) in degrees Centigrade
- Health status: healthy, nearing wearout, write-reduced or read-only
- Reserve capacity (percentage)
- Warning capacity threshold (percentage)

Data volume information: If the -d option is used, the following data volume information is reported in addition to the basic information:

- Physical bytes written
- Physical bytes read

All information: If the `-a` option is used, all information is printed, which includes the following information in addition to basic and data volume information.

Adapter information:

- Manufacturer number
- Part number
- Date of manufacture
- Power loss protection status
- PCIe bus voltage (average, minimum, maximum)
- PCIe bus current (average, maximum)
- PCIe bus power (average, maximum)
- PCIe power limit threshold (watts)
- PCIe slot available power (watts)
- PCIe negotiated link information (lanes and throughput)

Block device information:

- Manufacturer's code
- Manufacturing date
- Vendor and sub-vendor information
- Format status and sector information (if device is attached)
- FPGA ID and Low-level format GUID
- PCIe slot available power
- PCIe negotiated link information
- Card temperature, in degrees Centigrade
- Internal voltage, average and maximum.
- Auxiliary voltage: average and maximum.
- Percentage of good blocks, data and metadata
- Lifetime data volume statistics
- RAM usage

Error mode information: If the IO Accelerator is in minimal mode, read-only mode, or write-reduced mode when `fio-status` is run, the following differences occur in the output:

- Attach status is Status unknown: Driver is in MINIMAL MODE:
- The reason for the minimal mode state appears (for example, Firmware is out of date. Update firmware.)
- "Geometry and capacity information not available." is displayed.
- No media health information appears.

fio-sure-erase

CAUTION: Do not use this utility if there are any IO Accelerator devices installed in the system that are not selected to be cleared or purged of data.

- Ensure that you back up any data prior to activation of this utility.
- Remove any devices that are not targeted for purge.
- After the data is removed from the target devices, it is purged.
- There is no recovery to this action.

CAUTION: If the device is in Read-only mode, perform a format using `fio-format` before running `fio-sure-erase`.
The `fio-sure-erase` utility cannot erase the device if it is in Minimal mode. Updating the firmware might move the device out of Minimal Mode. If the device remains in Minimal mode, contact HP Support (http://www.hp.com/support) for assistance.

IMPORTANT: Prior to reactivating the device, format the device with `fio-format` after running `fio-sure-erase`.

To run `fio-sure-erase`, the block device must be detached. For more information, see "`fio-detach` (on page 24)".

Description

The `fio-sure-erase` is a command-line utility that securely removes data from IO Accelerator devices. It complies with the Clear and Purge level of destruction from the following standards:

- DOD 5220.22-M—Complies with instructions for Flash EPROM.
- NIST SP800-88—Complies with instructions for Flash EPROM.

For more information, see the following sections on "Clear support" and "Purge support."

Syntax

```
fio-sure-erase [options] <device>
```

where `<device>` is the name of the device node (`/dev/fctx`), where `x` indicates the card number: `0`, `1`, `2`, and so on. For example, `/dev/fct0` indicates the first IO Accelerator device installed on the system. To view this device node, use "`fio-status` (on page 26)".

IMPORTANT: **Products with Multiple Devices**
`fio-sure-erase` applies to individual IO Accelerator devices. For example, if you are planning to purge an ioDrive Duo device, perform this operation on each of the two IO Accelerator devices.

Options	Description
`-p`	Purge instead of Clear: performs a write followed by an erase. For more information, see "Purge support."
`-y`	No confirmation: Does not require a yes/no response to execute the utility.
`-q`	Quiet: Does not display the status bar.

IMPORTANT: If `fio-sure-erase` is run without options, then a Clear is performed. For more information, see "Clear support."

Each block of memory consists of uniform 1 bits or 0 bits.

Clear support

A Clear is the default state of running `fio-sure-erase` (with no options), and refers to the act of performing a full low-level erase (every cell pushed to 1) of the entire NAND media, including retired erase blocks.

Metadata that is required for operation will not be destroyed (media event log, erase counts, physical bytes read/written, performance and thermal history), but any user-specific metadata will be destroyed.

The following describes the steps taken in the Clear operation:

1. Creates a unity map of every addressable block (this allows `fio-sure-erase` to address every block, including previously unmapped bad blocks).
2. For each block, performs an erase cycle (every cell is pushed to 1).
3. Restores the bad block map.

Formats the device (the purpose of this is to make the device usable again, the utility erases all of the headers during the clear).

Purge support

A Purge is implemented by using the `-p` option with `fio-sure-erase`. Purge refers to the act of first overwriting the entire NAND media (including retired erase blocks) with a single character (every cell written to logical 0), and then performing a full chip erase (every cell pushed to 1) across all media (including retired erase blocks).

Metadata that is required for operation will not be destroyed (media event log, erase counts, physical bytes read/written, performance and thermal history), but any user-specific metadata will be destroyed.

The Purge operation includes the following steps:

1. Creates a unity map of every addressable block (this allows `fio-sure-erase` to address every block, including previously unmapped bad blocks).
2. For each block, performs a write cycle (every cell is pushed to 0).
3. For each block, performs an erase cycle (every cell pushed to 1).
4. Restores the bad block map.

Formats the device (the purpose of this is to make the device usable again, the utility erases all of the headers during the clear).

fio-update-iodrive

CAUTION: HP strongly recommends that data is backed up on any IO Accelerator device before performing a firmware upgrade.

Description

Updates the IO Accelerator device's firmware. This utility scans the PCIe bus for all IO Accelerator devices and updates them.

A progress bar and percentage are shown for each device as the update completes.

CAUTION:

- During a firmware upgrade, it is critical to maintain steady power or risk failure of the IO Accelerator device. Connecting a qualified UPS is recommended prior to performing a firmware upgrade.
- It is critical to load the driver after each firmware upgrade step when scheduling sequential, multiple firmware upgrades (example: 1.2.7 to 2.1.0 to 2.3.1). If the driver is not loaded, the on-drive format will not be changed and there will be data loss.
- Data loss may occur if the IO Accelerator device firmware is downgraded. Contact HP Support (http://www.hp.com/support) for recommendations.
- The default to upgrade all IO Accelerator devices does not use the `-d` or `-s` option. The firmware is located in the `<ioaccelerator_version.fff>` file. Confirm that all devices need the firmware upgrade. The `-p` (Pretend) option, can be run to view the possible results of the update.
- Ensure that all IO Accelerator devices are detached before updating the firmware.
- **Upgrade Path**

  There is a specific upgrade path to follow when upgrading an IO Accelerator device. Consult the *Release Notes* for this IO Accelerator release before upgrading any IO Accelerator devices.

IMPORTANT: If you receive an error message when updating the firmware that instructs you to update the midprom information, contact HP Customer Support (http://www.hp.com/support).

To update one or more specific devices:

If the IO Accelerator is loaded, use the `-d` option with the device number.

Syntax

```
fio-update-iodrive [options] <iodrive_version.fff>
```

where `<iodrive_version.fff>` is the path and firmware archive file provided by HP. The default path is `/usr/share/fio/firmware`. This parameter is required.

Options	Description
`-d`	Updates the specified devices (by `fctx`, where `x` is the number of the device shown in `fio-status`). If this option is not specified, all devices are updated. Use the `-d` or `-s` options carefully. Updating the wrong IO Accelerator device could damage that device.
`-f`	Force upgrade (used primarily to downgrade to an earlier firmware version). If the IO Accelerator is not loaded, this option also requires the `-s` option. Use the `-f` option carefully. Updating the wrong IO Accelerator device could damage that device.
`-l`	List the firmware available in the archive.
`-p`	Pretend: Shows what updates would be done. However, the actual firmware is not modified.
`-c`	Clears locks placed on a device.

Options	Description
-q	Runs the update process without displaying the progress bar or percentage.
-y	Confirm all warning messages.
-s	Updates the devices in the specified slots using "*" as a wildcard for devices. The slots are identified in the following PCIe format (as shown in lspci): [[[[<domain>]:]<bus>]:][<slot>][.[<func>]]

All three external LED indicators light up during the update process.

Monitoring IO Accelerator health

NAND flash and component failure

The IO Accelerator is a highly fault-tolerant storage subsystem that provides many levels of protection against component failure and the loss nature of solid state storage. However, as in all storage subsystems, component failures might occur.

By pro-actively monitoring device age and health, you can ensure reliable performance over the intended product life.

Health metrics

The IO Accelerator manages block retirement using pre-determined retirement thresholds. The HP IO Accelerator Management Tool and the `fio-status` utilities show a health indicator that starts at 100 and counts down to 0. As certain thresholds are crossed, various actions are taken.

At the 10% healthy threshold, a one-time warning is issued. For more information, see "Health monitoring techniques."

At 0%, the device is considered unhealthy. It enters write-reduced mode, which somewhat prolongs its lifespan so data can be safely migrated off. In this state the IO Accelerator device behaves normally, except for the reduced write performance.

After the 0% threshold, the device will soon enter read-only mode, and any attempt to write to the IO Accelerator device causes an error. Some filesystems might require special mount options to mount a read-only block device in addition to specifying that the mount must be read-only.

For example, under Linux, `ext3` requires that `-o ro, noload` is used. The `noload` option tells the filesystem to not try and replay the journal.

Consider the read-only mode as a final opportunity to migrate data off the device, as device failure is more likely with continued use.

The IO Accelerator device might enter failure mode. In this case, the device is offline and inaccessible. This can be caused by an internal catastrophic failure, improper firmware upgrade procedures, or device wearout.

The IO Accelerator driver manages LEB retirement via use of pre-determined retirement thresholds. The IO Accelerator Management Tool and the `fio-status` utility show a health indicator that starts at 100 and counts down to 0. As certain thresholds are crossed, various actions are taken.

At the 10% healthy threshold, a one-time warning is issued. For more information, see "Health monitoring techniques."

At 0%, the device is considered unhealthy. It enters write-reduced mode, which somewhat prolongs its lifespan so data can be safely migrated. In this state, the IO Accelerator behaves normally except for the reduced write performance.

At some point after the 0% threshold, the device enters read-only mode. Any attempt to write to the IO Accelerator causes an error. Some file systems might require special mount options to mount a read-only

block device, beyond specifying that the mount must be read-only. For example, under Linux, `ext3` requires that `-o ro, noload` is used. The `noload` option directs the file system not to attempt to replay the journal.

Read-only mode must be considered a final opportunity to migrate data off the device since device failure is more likely with continued use.

The IO Accelerator might enter failure mode. In this case, the device is offline and inaccessible. This can be caused by an internal catastrophic failure, improper firmware upgrade procedures, or device wears out.

IMPORTANT:

- For service or warranty-related questions, contact the company you purchased the device from.
- For products with multiple IO Accelerator devices, these modes are maintained independently for each device.

Health monitoring techniques

fio-status

Output from the `fio-status` utility shows the health percentage and drive state. These items are referenced as `Media status` in the following sample output.

```
Found 3 ioMemory devices in this system
Fusion-io driver version: 3.1.0 build 364
Adapter: Single Adapter
        HP IO Accelerator 1.30TB, Product Number:AJ878B,
SN:1133D0248, FIO SN:1134D9565
...
Media status: Healthy; Reserves: 100.00%, warn at 10.00%; Data: 99.12%
Lifetime data volumes:
   Physical bytes written: 6,423,563,326,064
   Physical bytes read   : 5,509,006,756,312
```

HP IO Accelerator Management Tool: In the Device Report tab, look for the Reserve Space percentage in the right column. The higher the percentage, the healthier the drive is likely to be.

SNMP: On Windows or Linux operating systems, see the corresponding section for details on "Configuring the SNMP master agent."

The following Health Status messages are produced by the `fio-status` utility:

- Healthy
- Read-only
- Reduced-write
- Unknown

ZFS pools and health monitoring

ZFS pools are typically configured to detect and mitigate the failure modes of traditional storage media. The IO Accelerator device attempts to fail as gracefully as possible, and these new failure mechanisms are compatible with ZFS pools. A device in a zpool fails to receive data at a sufficient rate if the device is in a write-reduced state and it is participating in a write-heavy workload. In this case, the drive is evicted from the zpool. A drive in read-only mode is evicted when write I/Os are returned from the device as failed. Catastrophic failures are detected and handled just as though they are on traditional storage devices.

Using module parameters

The following table describes the module parameters you can set using the `esxcfg-module` command.

The remote option (`--server`) is only required for the vCLI.

Sample command:

```
esxcfg-module --server <server-name> iomemory-vsl -s '<parameter>=<value>'
```

You must reboot the ESX or ESXi system to enforce any parameter changes.

Module parameter	Default (min/max)	Description
`auto_attach`	1	1 = Always attach the device on startup. 0 = Do not attach the device on startup.
external_power_override*	No devices selected	Allows selected devices to draw full power from the PCIe slot. The `<value>` for this parameter is a comma-separated list of adapter serial numbers.
`force_minimal_mode`	0	1 = Force minimal mode on the device. 0 = Do not force minimal mode on the device.
`parallel_attach`	True	1 = Enable parallel attach of multiple drivers. 0 = Disable parallel attach of multiple drives.
`tintr_hw_wait`	0 (0, 255)	Interval (microseconds) to wait between hardware interrupts. Also known as interrupt coalescing. 0 = Off.

*Use this parameter with care. For more information, see "Enabling PCIe power override (on page 14)."

Except for the `external_power_override` parameter, module parameters apply to all IO Accelerator devices in the system.

To query the current module parameters, use the following command:

```
esxcfg-module --server <server-name> iomemory-vsl -g
```

Upgrading devices

Upgrading devices for IO Accelerator VSL 2.x.x to 3.x.x

This version of IO Accelerator VSL software supports new features, including the latest generation of IO Accelerator architecture and improved Flashback protection. These features require the latest version of the IO Accelerator firmware. Every device in a system running 2.x or later must be upgraded to the latest version of the firmware.

For example, if you have a system running 2.x.x IO Accelerator VSL software with devices previously installed and you want to install new devices that require the latest version of the firmware, then you must upgrade all of the existing devices to the latest firmware version.

△ **CAUTION:** You cannot revert device firmware to an earlier version after you have upgraded the device without voiding your warranty. If you experience problems with your upgrade, contact HP Customer Support.

△ **CAUTION:** Upgrading devices (previously configured for VSL 2.x.x) to work with VSL 3.x.x requires a low-level media format of the device. No user data is maintained during the process. Be sure to back up all data as instructed.

Multiple upgrade passes for some versions

Depending on the current firmware version of your devices, you may need to upgrade your device's firmware multiple times to preserve internal structures. See the IO Accelerator VSL software for the upgrade path.

Upgrade procedure

Be sure to follow the upgrade path and make sure that all previously installed IO Accelerator devices are updated with the appropriate 3.2.3-compatible firmware.

If you plan to use IO Accelerator Gen1 devices and IO Accelerator Gen2 devices in the same host, perform this upgrade on all existing IO Accelerator Gen1 devices before installing the new IO Accelerator Gen2 devices.

1. Prepare each existing IO Accelerator device for upgrade:
   a. Back up user data on each IO Accelerator device.

   △ **CAUTION:** Upgrading IO Accelerator devices that were previously configured for VSL 1.x.x or 2.x.x to work with VSL 3.x.x requires a low-level media format of the device. No user data is maintained during the media format process. Be sure to backup all data on your IO Accelerator device as instructed before upgrading the firmware.

   Do not back up the data onto another IO Accelerator device on the same system. The backup must be to a local disk or to an externally attached volume.

b. Run the `fio-bugreport` command-line utility and save the output. This output captures the device information for each device in the system. This device information will be useful in troubleshooting any upgrade issues. For example:

```
fio-bugreport
```

c. Detach the IO Accelerator devices. For example:

```
fio-detach /dev/fct*
```

For more information, see "`fio-detach` (on page 24)."

2. Unload the current IO Accelerator driver. For example:

```
$ pfexec rem_drv iomemory-vsl
```

3. Uninstall the 2.x HP IO Accelerator VSL software:

```
$ pfexec pkgrm iomemory-vsl
```

4. Install the new VSL and related packages:

a. Download the IO Accelerator VSL package and supporting packages from the HP website (http://www.hp.com/support).

b. Uncompress the package:

```
$ gunzip -c iomemory-vsl*pkg.tar.gz|tar -xf -
```

c. Install the VSL and utilities using the following command:

```
$ pfexec pkgadd -d . iomemory-vsl
```

5. Update the firmware on each device to the latest version using the `fio-update-iodrive` command.

CAUTION: Do not turn off the power during a firmware upgrade, because this might cause device failure. If a UPS is not in place, consider adding one to the system before performing a firmware upgrade.

Sample syntax:

```
fio-update-iodrive [options] <iodrive_version.fff>
```

Where `<iodrive_version.fff>` is the path to the firmware archive. This command updates all of the devices to the selected firmware. See the utility reference for more options If you want to update specific devices.

6. Reboot the system.

A warning that the upgraded devices are missing a lebmap, if `fio-status` is run. This is a customary warning, and the issue will be corrected in the next step.

7. Load the VSL. For example:

```
$ modprobe iomemory-vsl
```

For more information, see "Loading the IO Accelerator driver ("Loading the IO Accelerator VSL facility (driver)" on page 12)."

CAUTION: Use this utility with care since it deletes all user information on the IO Accelerator.

8. Using the `fio-format` command, format each device. For example:

```
fio-format <device>
```

You are prompted to confirm you want to erase all data on the device. The format might take an extended period of time, depending on the wear on the device.

9. Attach all IO Accelerator devices using the following command:

```
fio-attach /dev/fct*
```

10. Using the following command, check the status of all devices:

    ```
    fio-status -a
    ```

Your IO Accelerator devices are now successfully upgraded for this version of the HP IO Accelerator. You can now install any IO Accelerator Gen2 devices.

Resources

Subscription service

HP recommends that you register your product at the Subscriber's Choice for Business website (http://www.hp.com/support).

After registering, you will receive e-mail notification of product enhancements, new driver versions, firmware updates, and other product resources.

For more information

For additional information, see the following HP websites:

- HP BladeSystem technical resources (http://www.hp.com/go/bladesystem/documentation) (white papers and support documents)
- HP BladeSystem components (http://h18004.www1.hp.com/products/blades/components/c-class-compmatrix.html)
- HP support (http://www.hp.com/support)

NOTE: Before contacting HP customer support, run the IO Accelerator bug reporting tool, and have the report with you when you call. To run the IO Accelerator bug reporting tool, enter the `fio-bugreport` command.

Support and other resources

Before you contact HP

Be sure to have the following information available before you call HP:

- Active Health System log (HP ProLiant Gen8 or later products)

  Download and have available an Active Health System log for 3 days before the failure was detected. For more information, see the *HP iLO 4 User Guide* or *HP Intelligent Provisioning User Guide* on the HP website (http://www.hp.com/go/ilo/docs).
- Onboard Administrator SHOW ALL report (for HP BladeSystem products only)

  For more information on obtaining the Onboard Administrator SHOW ALL report, see the HP website (http://www.hp.com/go/OAlog).
- Technical support registration number (if applicable)
- Product serial number
- Product model name and number
- Product identification number
- Applicable error messages
- Add-on boards or hardware
- Third-party hardware or software
- Operating system type and revision level

HP contact information

For United States and worldwide contact information, see the Contact HP website (http://www.hp.com/go/assistance).

In the United States:

- To contact HP by phone, call 1-800-334-5144. For continuous quality improvement, calls may be recorded or monitored.
- If you have purchased a Care Pack (service upgrade), see the Support & Drivers website (http://www8.hp.com/us/en/support-drivers.html). If the problem cannot be resolved at the website, call 1-800-633-3600. For more information about Care Packs, see the HP website (http://pro-aq-sama.houston.hp.com/services/cache/10950-0-0-225-121.html).

Customer Self Repair

HP products are designed with many Customer Self Repair (CSR) parts to minimize repair time and allow for greater flexibility in performing defective parts replacement. If during the diagnosis period HP (or HP service

providers or service partners) identifies that the repair can be accomplished by the use of a CSR part, HP will ship that part directly to you for replacement. There are two categories of CSR parts:

- **Mandatory**—Parts for which customer self repair is mandatory. If you request HP to replace these parts, you will be charged for the travel and labor costs of this service.
- **Optional**—Parts for which customer self repair is optional. These parts are also designed for customer self repair. If, however, you require that HP replace them for you, there may or may not be additional charges, depending on the type of warranty service designated for your product.

NOTE: Some HP parts are not designed for customer self repair. In order to satisfy the customer warranty, HP requires that an authorized service provider replace the part. These parts are identified as "No" in the Illustrated Parts Catalog.

Based on availability and where geography permits, CSR parts will be shipped for next business day delivery. Same day or four-hour delivery may be offered at an additional charge where geography permits. If assistance is required, you can call the HP Technical Support Center and a technician will help you over the telephone. HP specifies in the materials shipped with a replacement CSR part whether a defective part must be returned to HP. In cases where it is required to return the defective part to HP, you must ship the defective part back to HP within a defined period of time, normally five (5) business days. The defective part must be returned with the associated documentation in the provided shipping material. Failure to return the defective part may result in HP billing you for the replacement. With a customer self repair, HP will pay all shipping and part return costs and determine the courier/carrier to be used.

For more information about HP's Customer Self Repair program, contact your local service provider. For the North American program, refer to the HP website (http://www.hp.com/go/selfrepair).

Réparation par le client (CSR)

Les produits HP comportent de nombreuses pièces CSR (Customer Self Repair = réparation par le client) afin de minimiser les délais de réparation et faciliter le remplacement des pièces défectueuses. Si pendant la période de diagnostic, HP (ou ses partenaires ou mainteneurs agréés) détermine que la réparation peut être effectuée à l'aide d'une pièce CSR, HP vous l'envoie directement. Il existe deux catégories de pièces CSR:

Obligatoire - Pièces pour lesquelles la réparation par le client est obligatoire. Si vous demandez à HP de remplacer ces pièces, les coûts de déplacement et main d'œuvre du service vous seront facturés.

Facultatif - Pièces pour lesquelles la réparation par le client est facultative. Ces pièces sont également conçues pour permettre au client d'effectuer lui-même la réparation. Toutefois, si vous demandez à HP de remplacer ces pièces, l'intervention peut ou non vous être facturée, selon le type de garantie applicable à votre produit.

REMARQUE: Certaines pièces HP ne sont pas conçues pour permettre au client d'effectuer lui-même la réparation. Pour que la garantie puisse s'appliquer, HP exige que le remplacement de la pièce soit effectué par un Mainteneur Agréé. Ces pièces sont identifiées par la mention "Non" dans le Catalogue illustré.

Les pièces CSR sont livrées le jour ouvré suivant, dans la limite des stocks disponibles et selon votre situation géographique. Si votre situation géographique le permet et que vous demandez une livraison le jour même ou dans les 4 heures, celle-ci vous sera facturée. Pour bénéficier d'une assistance téléphonique, appelez le Centre d'assistance technique HP. Dans les documents envoyés avec la pièce de rechange CSR, HP précise s'il est nécessaire de lui retourner la pièce défectueuse. Si c'est le cas, vous devez le faire dans le délai indiqué, généralement cinq (5) jours ouvrés. La pièce et sa documentation doivent être retournées dans l'emballage fourni. Si vous ne retournez pas la pièce défectueuse, HP se réserve le droit de vous facturer les coûts de remplacement. Dans le cas d'une pièce CSR, HP supporte l'ensemble des frais d'expédition et de retour, et détermine la société de courses ou le transporteur à utiliser.

Pour plus d'informations sur le programme CSR de HP, contactez votre Mainteneur Agrée local. Pour plus d'informations sur ce programme en Amérique du Nord, consultez le site Web HP (http://www.hp.com/go/selfrepair).

Riparazione da parte del cliente

Per abbreviare i tempi di riparazione e garantire una maggiore flessibilità nella sostituzione di parti difettose, i prodotti HP sono realizzati con numerosi componenti che possono essere riparati direttamente dal cliente (CSR, Customer Self Repair). Se in fase di diagnostica HP (o un centro di servizi o di assistenza HP) identifica il guasto come riparabile mediante un ricambio CSR, HP lo spedirà direttamente al cliente per la sostituzione. Vi sono due categorie di parti CSR:

Obbligatorie – Parti che devono essere necessariamente riparate dal cliente. Se il cliente ne affida la riparazione ad HP, deve sostenere le spese di spedizione e di manodopera per il servizio.

Opzionali – Parti la cui riparazione da parte del cliente è facoltativa. Si tratta comunque di componenti progettati per questo scopo. Se tuttavia il cliente ne richiede la sostituzione ad HP, potrebbe dover sostenere spese addizionali a seconda del tipo di garanzia previsto per il prodotto.

NOTA: alcuni componenti HP non sono progettati per la riparazione da parte del cliente. Per rispettare la garanzia, HP richiede che queste parti siano sostituite da un centro di assistenza autorizzato. Tali parti sono identificate da un "No" nel Catalogo illustrato dei componenti.

In base alla disponibilità e alla località geografica, le parti CSR vengono spedite con consegna entro il giorno lavorativo seguente. La consegna nel giorno stesso o entro quattro ore è offerta con un supplemento di costo solo in alcune zone. In caso di necessità si può richiedere l'assistenza telefonica di un addetto del centro di supporto tecnico HP. Nel materiale fornito con una parte di ricambio CSR, HP specifica se il cliente deve restituire dei componenti. Qualora sia richiesta la resa ad HP del componente difettoso, lo si deve spedire ad HP entro un determinato periodo di tempo, generalmente cinque (5) giorni lavorativi. Il componente difettoso deve essere restituito con la documentazione associata nell'imballo di spedizione fornito. La mancata restituzione del componente può comportare la fatturazione del ricambio da parte di HP. Nel caso di riparazione da parte del cliente, HP sostiene tutte le spese di spedizione e resa e sceglie il corriere/vettore da utilizzare.

Per ulteriori informazioni sul programma CSR di HP contattare il centro di assistenza di zona. Per il programma in Nord America fare riferimento al sito Web HP (http://www.hp.com/go/selfrepair).

Customer Self Repair

HP Produkte enthalten viele CSR-Teile (Customer Self Repair), um Reparaturzeiten zu minimieren und höhere Flexibilität beim Austausch defekter Bauteile zu ermöglichen. Wenn HP (oder ein HP Servicepartner) bei der Diagnose feststellt, dass das Produkt mithilfe eines CSR-Teils repariert werden kann, sendet Ihnen HP dieses Bauteil zum Austausch direkt zu. CSR-Teile werden in zwei Kategorien unterteilt:

Zwingend – Teile, für die das Customer Self Repair-Verfahren zwingend vorgegeben ist. Wenn Sie den Austausch dieser Teile von HP vornehmen lassen, werden Ihnen die Anfahrt- und Arbeitskosten für diesen Service berechnet.

Optional – Teile, für die das Customer Self Repair-Verfahren optional ist. Diese Teile sind auch für Customer Self Repair ausgelegt. Wenn Sie jedoch den Austausch dieser Teile von HP vornehmen lassen möchten, können bei diesem Service je nach den für Ihr Produkt vorgesehenen Garantiebedingungen zusätzliche Kosten anfallen.

HINWEIS: Einige Teile sind nicht für Customer Self Repair ausgelegt. Um den Garantieanspruch des Kunden zu erfüllen, muss das Teil von einem HP Servicepartner ersetzt werden. Im illustrierten Teilekatalog sind diese Teile mit „No“ bzw. „Nein“ gekennzeichnet.

CSR-Teile werden abhängig von der Verfügbarkeit und vom Lieferziel am folgenden Geschäftstag geliefert. Für bestimmte Standorte ist eine Lieferung am selben Tag oder innerhalb von vier Stunden gegen einen Aufpreis verfügbar. Wenn Sie Hilfe benötigen, können Sie das HP technische Support Center anrufen und sich von einem Mitarbeiter per Telefon helfen lassen. Den Materialien, die mit einem CSR-Ersatzteil geliefert werden, können Sie entnehmen, ob das defekte Teil an HP zurückgeschickt werden muss. Wenn es erforderlich ist, das defekte Teil an HP zurückzuschicken, müssen Sie dies innerhalb eines vorgegebenen Zeitraums tun, in der Regel innerhalb von fünf (5) Geschäftstagen. Das defekte Teil muss mit der zugehörigen Dokumentation in der Verpackung zurückgeschickt werden, die im Lieferumfang enthalten ist. Wenn Sie das defekte Teil nicht zurückschicken, kann HP Ihnen das Ersatzteil in Rechnung stellen. Im Falle von Customer Self Repair kommt HP für alle Kosten für die Lieferung und Rücksendung auf und bestimmt den Kurier-/Frachtdienst.

Weitere Informationen über das HP Customer Self Repair Programm erhalten Sie von Ihrem Servicepartner vor Ort. Informationen über das CSR-Programm in Nordamerika finden Sie auf der HP Website unter (http://www.hp.com/go/selfrepair).

Reparaciones del propio cliente

Los productos de HP incluyen muchos componentes que el propio usuario puede reemplazar (*Customer Self Repair*, CSR) para minimizar el tiempo de reparación y ofrecer una mayor flexibilidad a la hora de realizar sustituciones de componentes defectuosos. Si, durante la fase de diagnóstico, HP (o los proveedores o socios de servicio de HP) identifica que una reparación puede llevarse a cabo mediante el uso de un componente CSR, HP le enviará dicho componente directamente para que realice su sustitución. Los componentes CSR se clasifican en dos categorías:

- **Obligatorio:** componentes para los que la reparación por parte del usuario es obligatoria. Si solicita a HP que realice la sustitución de estos componentes, tendrá que hacerse cargo de los gastos de desplazamiento y de mano de obra de dicho servicio.
- **Opcional:** componentes para los que la reparación por parte del usuario es opcional. Estos componentes también están diseñados para que puedan ser reparados por el usuario. Sin embargo, si precisa que HP realice su sustitución, puede o no conllevar costes adicionales, dependiendo del tipo de servicio de garantía correspondiente al producto.

NOTA: Algunos componentes no están diseñados para que puedan ser reparados por el usuario. Para que el usuario haga valer su garantía, HP pone como condición que un proveedor de servicios autorizado realice la sustitución de estos componentes. Dichos componentes se identifican con la palabra "No" en el catálogo ilustrado de componentes.

Según la disponibilidad y la situación geográfica, los componentes CSR se enviarán para que lleguen a su destino al siguiente día laborable. Si la situación geográfica lo permite, se puede solicitar la entrega en el mismo día o en cuatro horas con un coste adicional. Si precisa asistencia técnica, puede llamar al Centro de asistencia técnica de HP y recibirá ayuda telefónica por parte de un técnico. Con el envío de materiales para la sustitución de componentes CSR, HP especificará si los componentes defectuosos deberán devolverse a HP. En aquellos casos en los que sea necesario devolver algún componente a HP, deberá hacerlo en el periodo de tiempo especificado, normalmente cinco días laborables. Los componentes defectuosos deberán devolverse con toda la documentación relacionada y con el embalaje de envío. Si no enviara el componente defectuoso requerido, HP podrá cobrarle por el de sustitución. En el caso de todas

sustituciones que lleve a cabo el cliente, HP se hará cargo de todos los gastos de envío y devolución de componentes y escogerá la empresa de transporte que se utilice para dicho servicio.

Para obtener más información acerca del programa de Reparaciones del propio cliente de HP, póngase en contacto con su proveedor de servicios local. Si está interesado en el programa para Norteamérica, visite la página web de HP siguiente (http://www.hp.com/go/selfrepair).

Customer Self Repair

Veel onderdelen in HP producten zijn door de klant zelf te repareren, waardoor de reparatieduur tot een minimum beperkt kan blijven en de flexibiliteit in het vervangen van defecte onderdelen groter is. Deze onderdelen worden CSR-onderdelen (Customer Self Repair) genoemd. Als HP (of een HP Service Partner) bij de diagnose vaststelt dat de reparatie kan worden uitgevoerd met een CSR-onderdeel, verzendt HP dat onderdeel rechtstreeks naar u, zodat u het defecte onderdeel daarmee kunt vervangen. Er zijn twee categorieën CSR-onderdelen:

Verplicht: Onderdelen waarvoor reparatie door de klant verplicht is. Als u HP verzoekt deze onderdelen voor u te vervangen, worden u voor deze service reiskosten en arbeidsloon in rekening gebracht.

Optioneel: Onderdelen waarvoor reparatie door de klant optioneel is. Ook deze onderdelen zijn ontworpen voor reparatie door de klant. Als u echter HP verzoekt deze onderdelen voor u te vervangen, kunnen daarvoor extra kosten in rekening worden gebracht, afhankelijk van het type garantieservice voor het product.

OPMERKING: Sommige HP onderdelen zijn niet ontwikkeld voor reparatie door de klant. In verband met de garantievoorwaarden moet het onderdeel door een geautoriseerde Service Partner worden vervangen. Deze onderdelen worden in de geïllustreerde onderdelencatalogus aangemerkt met "Nee".

Afhankelijk van de leverbaarheid en de locatie worden CSR-onderdelen verzonden voor levering op de eerstvolgende werkdag. Levering op dezelfde dag of binnen vier uur kan tegen meerkosten worden aangeboden, indien dit mogelijk is gezien de locatie. Indien assistentie gewenst is, belt u een HP Service Partner om via de telefoon technische ondersteuning te ontvangen. HP vermeldt in de documentatie bij het vervangende CSR-onderdeel of het defecte onderdeel aan HP moet worden geretourneerd. Als het defecte onderdeel aan HP moet worden teruggezonden, moet u het defecte onderdeel binnen een bepaalde periode, gewoonlijk vijf (5) werkdagen, retourneren aan HP. Het defecte onderdeel moet met de bijbehorende documentatie worden geretourneerd in het meegeleverde verpakkingsmateriaal. Als u het defecte onderdeel niet terugzendt, kan HP u voor het vervangende onderdeel kosten in rekening brengen. Bij reparatie door de klant betaalt HP alle verzendkosten voor het vervangende en geretourneerde onderdeel en kiest HP zelf welke koerier/transportonderneming hiervoor wordt gebruikt.

Neem contact op met een Service Partner voor meer informatie over het Customer Self Repair programma van HP. Informatie over Service Partners vindt u op de HP website (http://www.hp.com/go/selfrepair).

Reparo feito pelo cliente

Os produtos da HP são projetados com muitas peças para reparo feito pelo cliente (CSR) de modo a minimizar o tempo de reparo e permitir maior flexibilidade na substituição de peças com defeito. Se, durante o período de diagnóstico, a HP (ou fornecedores/parceiros de serviço da HP) concluir que o reparo pode ser efetuado pelo uso de uma peça CSR, a peça de reposição será enviada diretamente ao cliente. Existem duas categorias de peças CSR:

Obrigatória – Peças cujo reparo feito pelo cliente é obrigatório. Se desejar que a HP substitua essas peças, serão cobradas as despesas de transporte e mão-de-obra do serviço.

Opcional – Peças cujo reparo feito pelo cliente é opcional. Essas peças também são projetadas para o reparo feito pelo cliente. No entanto, se desejar que a HP as substitua, pode haver ou não a cobrança de taxa adicional, dependendo do tipo de serviço de garantia destinado ao produto.

OBSERVAÇÃO: Algumas peças da HP não são projetadas para o reparo feito pelo cliente. A fim de cumprir a garantia do cliente, a HP exige que um técnico autorizado substitua a peça. Essas peças estão identificadas com a marca "No" (Não), no catálogo de peças ilustrado.

Conforme a disponibilidade e o local geográfico, as peças CSR serão enviadas no primeiro dia útil após o pedido. Onde as condições geográficas permitirem, a entrega no mesmo dia ou em quatro horas pode ser feita mediante uma taxa adicional. Se precisar de auxílio, entre em contato com o Centro de suporte técnico da HP para que um técnico o ajude por telefone. A HP especifica nos materiais fornecidos com a peça CSR de reposição se a peça com defeito deve ser devolvida à HP. Nos casos em que isso for necessário, é preciso enviar a peça com defeito à HP dentro do período determinado, normalmente cinco (5) dias úteis. A peça com defeito deve ser enviada com a documentação correspondente no material de transporte fornecido. Caso não o faça, a HP poderá cobrar a reposição. Para as peças de reparo feito pelo cliente, a HP paga todas as despesas de transporte e de devolução da peça e determina a transportadora/serviço postal a ser utilizado.

Para obter mais informações sobre o programa de reparo feito pelo cliente da HP, entre em contato com o fornecedor de serviços local. Para o programa norte-americano, visite o site da HP (http://www.hp.com/go/selfrepair).

カスタマーセルフリペア

修理時間を短縮し、故障部品の交換における高い柔軟性を確保するために、HP製品には多数のCSR部品があります。診断の際に、CSR部品を使用すれば修理ができるとHP（HPまたはHP正規保守代理店）が判断した場合、HPはその部品を直接、お客様に発送し、お客様に交換していただきます。CSR部品には以下の2通りがあります。

- **必須 -** カスタマーセルフリペアが必須の部品。当該部品について、もしもお客様がHPに交換作業を依頼される場合には、その修理サービスに関する交通費および人件費がお客様に請求されます。
- **任意 -** カスタマーセルフリペアが任意である部品。この部品もカスタマーセルフリペア用です。当該部品について、もしもお客様がHPに交換作業を依頼される場合には、お買い上げの製品に適用される保証サービス内容の範囲内においては、別途費用を負担していただくことなく保証サービスを受けることができます。

注： HP製品の一部の部品は、カスタマーセルフリペア用ではありません。製品の保証を継続するためには、HPまたはHP正規保守代理店による交換作業が必須となります。部品カタログには、当該部品がカスタマーセルフリペア除外品である旨が記載されています。

部品供給が可能な場合、地域によっては、CSR部品を翌営業日に届くように発送します。また、地域によっては、追加費用を負担いただくことにより同日または4時間以内に届くように発送することも可能な場合があります。サポートが必要なときは、HPの修理受付窓口に電話していただければ、技術者が電話でアドバイスします。交換用のCSR部品または同梱物には、故障部品をHPに返送する必要があるかどうかが表示されています。故障部品をHPに返送する必要がある場合は、指定期限内（通常は5営業日以内）に故障部品をHPに返送してください。故障部品を返送する場合は、届いた時の梱包箱に関連書類とともに入れてください。故障部品を返送しない場合、HPから部品費用が請求されます。カスタマーセルフリペアの際には、HPは送料および部品返送費を全額負担し、使用する宅配便会社や運送会社を指定します。

客户自行维修

HP 产品提供许多客户自行维修 (CSR) 部件，以尽可能缩短维修时间和在更换缺陷部件方面提供更大的灵活性。如果在诊断期间 HP（或 HP 服务提供商或服务合作伙伴）确定可以通过使用 CSR 部件完成维修，HP 将直接把该部件发送给您进行更换。有两类 CSR 部件：

- **强制性的** — 要求客户必须自行维修的部件。如果您请求 HP 更换这些部件，则必须为该服务支付差旅费和人工费用。
- **可选的** — 客户可以选择是否自行维修的部件。这些部件也是为客户自行维修设计的。不过，如果您要求 HP 为您更换这些部件，则根据为您的产品指定的保修服务类型，HP 可能收取或不再收取任何附加费用。

注：某些 HP 部件的设计并未考虑客户自行维修。为了满足客户保修的需要，HP 要求授权服务提供商更换相关部件。这些部件在部件图解目录中标记为“否”。

CSR 部件将在下一个工作日发运（取决于备货情况和允许的地理范围）。在允许的地理范围内，可在当天或四小时内发运，但要收取额外费用。如果需要帮助，您可以致电 HP 技术支持中心，将会有技术人员通过电话为您提供帮助。HP 会在随更换的 CSR 部件发运的材料中指明是否必须将有缺陷的部件返还给 HP。如果要求您将有缺陷的部件返还给 HP，那么您必须在规定期限内（通常是五 (5) 个工作日）将缺陷部件发给 HP。有缺陷的部件必须随所提供的发运材料中的相关文件一起返还。如果未能送还有缺陷的部件，HP 可能会要求您支付更换费用。客户自行维修时，HP 将承担所有相关运输和部件返回费用，并指定快递商/承运商。

有关 HP 客户自行维修计划的详细信息，请与您当地的服务提供商联系。有关北美地区的计划，请访问 HP 网站 (http://www.hp.com/go/selfrepair)。

客戶自行維修

HP 產品設計了許多「客戶自行維修」(CSR) 的零件以減少維修時間，並且使得更換瑕疵零件時能有更大的彈性。如果在診斷期間 HP（或 HP 服務供應商或維修夥伴）辨認出此項維修工作可以藉由使用 CSR 零件來完成，則 HP 將直接寄送該零件給您作更換。CSR 零件分為兩種類別：

- **強制的** — 客戶自行維修所使用的零件是強制性的。如果您要求 HP 更換這些零件，HP 將會向您收取此服務所需的外出費用與勞動成本。
- **選購的** — 客戶自行維修所使用的零件是選購的。這些零件也設計用於客戶自行維修之用。不過，如果您要求 HP 為您更換，則可能需要也可能不需要負擔額外的費用，端視針對此產品指定的保固服務類型而定。

備註：某些 HP 零件沒有消費者可自行維修的設計。為符合客戶保固，HP 需要授權的服務供應商更換零件。這些零件在圖示的零件目錄中，被標示為「否」。

基於材料取得及環境允許的情況下，CSR 零件將於下一個工作日以快遞寄送。在環境的允許下當天或四小時內送達，則可能需要額外的費用。若您需要協助，可致電「HP 技術支援中心」，會有一位技術人員透過電話來協助您。不論損壞的零件是否必須退回，HP 皆會在與 CSR 替換零件一起運送的材料中註明。若要將損壞的零件退回 HP，您必須在指定的一段時間內（通常為五 (5) 個工作天），將損壞的零件寄回 HP。損壞的零件必須與寄送資料中隨附的相關技術文件一併退還。如果無法退還損壞的零件，HP 可能要向您收取替換費用。針對客戶自行維修情形，HP 將負責所有運費及零件退還費用並指定使用何家快遞/貨運公司。

如需 HP 的「客戶自行維修」方案詳細資訊，請連絡您當地的服務供應商。至於北美方案，請參閱 HP 網站 (http://www.hp.com/go/selfrepair)。

고객 셀프 수리

HP 제품은 수리 시간을 최소화하고 결함이 있는 부품 교체 시 더욱 융통성을 발휘할 수 있도록 하기 위해 고객 셀프 수리(CSR) 부품을 다량 사용하여 설계되었습니다. 진단 기간 동안 HP(또는 HP 서비스 공급업체 또는 서비스 협력업체)에서 CSR 부품을 사용하여 수리가 가능하다고 판단되면 HP는 해당 부품을 바로 사용자에게 보내어 사용자가 교체할 수 있도록 합니다. CSR 부품에는 두 가지 종류가 있습니다.

- **고객 셀프 수리가 의무 사항인 필수 부품.** 사용자가 HP에 이 부품의 교체를 요청할 경우 이 서비스에 대한 출장비 및 작업비가 청구됩니다.
- **고객 셀프 수리가 선택 사항인 부품.** 이 부품들도 고객 셀프 수리가 가능하도록 설계되었습니다. 하지만 사용자가 HP에 이 부품의 교체를 요청할 경우 사용자가 구입한 제품에 해당하는 보증 서비스 유형에 따라 추가 비용 없이 교체가 가능할 수 있습니다.

참고: 일부 HP 부품은 고객 셀프 수리가 불가능하도록 설계되었습니다. HP는 만족스러운 고객 보증을 위해 공인 서비스 제공업체를 통해 부품을 교체하도록 하고 있습니다. 이러한 부품들은 Illustrated Parts Catalog에 "No"라고 표시되어 있습니다.

CSR 부품은 재고 상태와 지리적 조건이 허용하는 경우 다음 영업일 납품이 가능하도록 배송이 이루어집니다. 지리적 조건이 허용하는 경우 추가 비용이 청구되는 조건으로 당일 또는 4시간 배송이 가능할 수도 있습니다. 도움이 필요하시면 HP 기술 지원 센터로 전화하십시오. 전문 기술자가 전화로 도움을 줄 것입니다. HP는 결함이 발생한 부품을 HP로 반환해야 하는지 여부를 CSR 교체 부품과 함께 배송된 자료에 지정합니다. 결함이 발생한 부품을 HP로 반환해야 하는 경우에는 지정된 기간 내(통상 영업일 기준 5일)에 HP로 반환해야 합니다. 이 때 결함이 발생한 부품은 제공된 포장 재료에 넣어 관련 설명서와 함께 반환해야 합니다. 결함이 발생한 부품을 반환하지 않는 경우 HP가 교체 부품에 대해 비용을 청구할 수 있습니다. 고객 셀프 수리의 경우, HP는 모든 운송 및 부품 반환 비용을 부담하며 이용할 운송업체 및 택배 서비스를 결정합니다.

HP 고객 셀프 수리 프로그램에 대한 자세한 내용은 가까운 서비스 제공업체에 문의하십시오. 북미 지역의 프로그램에 대해서는 HP 웹 사이트(http://www.hp.com/go/selfrepair)를 참조하십시오.

Acronyms and abbreviations

ACPI

Advanced Configuration and Power Interface

DVFS

dynamic voltage and frequency scaling

IOPS

input/output operations per second

LEB

Logical Erase Block

LVM

Logical Volume Manager

MIB

management information base

NAND

Not AND

RHEL

Red Hat Enterprise Linux

RPM

Red Hat Package Manager

VSL

virtual storage layer

Documentation feedback

HP is committed to providing documentation that meets your needs. To help us improve the documentation, send any errors, suggestions, or comments to Documentation Feedback (mailto:docsfeedback@hp.com). Include the document title and part number, version number, or the URL when submitting your feedback.

Index

A

about this guide 5
authorized reseller 43, 44

C

command-line utilities 15
configuring a ZFS pool 13
contents summary 5
CSR (customer self repair) 44
customer self repair (CSR) 44

D

devices, upgrading 39
disabling auto attach 17
disabling DVFS 23
disabling the driver 18
documentation feedback 53

E

Enabling PCIe power 13
enabling PCIe power override 15
enabling the driver 18
enabling the override parameter 16

F

fio-attach utility 25
fio-beacon utility 26
fio-bugreport utility 26
fio-detach utility 27
fio-format utility 28
fio-status utility 29
fio-sure-erase utility 31
fio-update-iodrive utility 33
firmware, upgrading 12
For more information 42

H

health metrics 36
health monitoring 38
health monitoring techniques 37
HP Subscriber's Choice for Business 42
HP technical support 43, 44
HP, contacting 44

I

installation overview 7, 11
installing software 11
introduction 6

L

LED indicators 22
limiting APCI C-states 23
loading the driver 12

M

maintenance 15, 18
maintenance tools 15
management tools 19
managing devices 19
module parameters 38
monitoring conditions 20
monitoring devices 19
monitoring IO Accelerator health 36

N

NAND flash and component failure 36

O

operating environment 8
overview 6

P

PCIe power override, enabling 15
PCIe power, enabling 13
performance and tuning 23
performance and tuning, introduction 23
performance attributes 7
phone numbers 44
product naming 6

R

resources 42, 43

S

software installation 11
support 43, 44
support and other resources 43
supported firmware revisions 8
supported hardware 9

T

technical support 43, 44
telephone numbers 43, 44

U

uninstalling the driver 17
unmanaged shutdown issues 18
upgrading devices 39
upgrading firmware 12
upgrading, procedure 39
using module parameters 38
using the IO Accelerator as a swap 13
utilities 11, 25
utilities reference 25

W

website, HP 44

Z

ZFS pool, configuring 13
ZFS pools 38